

## **REMARKS**

### **I. General**

The present application has 28 claims pending, all of which stand rejected in the current Office Action mailed June 20, 2006. The issues raised in the current Office Action are:

- Claims 1-28 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1-47 of copending Application 10/345,587 in view of U.S. Patent Publication 2004/0088380 to Chung et al. (hereinafter "*Chung*").
- Claims 1-28 are rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent 6,970,939 to Siew Yong Sim (hereinafter "*Sim*") in view of U.S. Patent 6,477,583 to Zayas et al. (hereinafter "*Zayas*"), and further in view of *Chung*.

Applicant hereby traverses all outstanding rejections and respectfully requests reconsideration and withdrawal thereof in light of the following remarks.

### **II. Amendment**

Claim 20 is amended herein to correct a typographical error by removing a redundant "said". This is intended solely as a cosmetic change to the claim, and is not intended to narrow its scope in any way.

### **III. The Obviousness-Type Double Patent Rejection**

Claims 1-28 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting over claims 1-47 of copending Application 10/345,587 ("the '587 application") in view of *Chung*. Applicant respectfully traverses for the reasons stated below.

An obviousness-type double patenting rejection should make clear the differences between the inventions defined by the conflicting claims—a claim in the patent compared to a claim in the application, and the reasons why a person of ordinary skill in the art would conclude that the invention defined in the claim at issue would have been an obvious variation of the invention defined in a claim in the patent. M.P.E.P. § 804(II)(B)(1). Thus, to establish a proper

grounds of obviousness-type double patenting, the Examiner must apply the *Graham* obviousness analysis. The present Office Action fails to properly establish such obviousness-type double patenting for the reasons discussed below.

The Office Action concedes at pages 2-3 thereof that claims 1, 22, and 36 of the '587 application fails to recite that at least one recipient node of a first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of the first group before the at least one recipient node fully receives its respective subfile. However, the Office Action contends that *Chung* teaches this element. Applicant respectfully disagrees. For the reasons discussed further herein below, *Chung* does not teach or suggest this element, and thus the Office Action fails to establish a proper obviousness-type double patenting rejection. Further, the Office Action fails to establish any reasoning why one of ordinary skill in the art would have been motivated to combine the teaching of *Chung* with the claims of the '587 application.

Accordingly, Applicant respectfully requests that the Examiner withdraw the obviousness-type double patenting rejection of record.

#### **IV. The 35 U.S.C. §103 Rejections over *Sim* in view of *Zayas* and further in view of *Chung***

Claims 1-28 are rejected under 35 U.S.C. §103(a) as being unpatentable over *Sim* in view of *Zayas* and further in view of *Chung*. Applicant respectfully traverses these rejections below.

To establish a prima facie case of obviousness, three basic criteria must be met. *See* M.P.E.P. § 2143. First, there must be some suggestion or motivation, either in the applied references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the applied references must teach or suggest all the claim limitations. Without conceding any other criteria, Applicant respectfully asserts that the combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all of the claim limitations, and

the Examiner has failed to properly establish sufficient motivation for combining the *Sim*, *Zayas*, and *Chung* references, as discussed further below.

*1. The Applied Combination Fails to Teach or Suggest All Claim Elements*

Independent Claim 1

Independent claim 1 recites:

A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:  
partitioning a file *F* into a plurality of subfiles;  
performing distribution of said file *F* to a plurality of recipient nodes using a distribution technique that comprises  
(a) attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of said first group but not all of said plurality of subfiles to any recipient node of said first group, and  
(b) said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;  
detecting a failed node of said plurality of recipient nodes; and  
said distribution technique adapting to distribute all of the subfiles of said file *F* to each non-failed node of said plurality of recipient nodes. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 1. As discussed below, the applied combination fails to teach or suggest at least: 1) a plurality of recipient nodes attempting to exchange their respective subfiles received from a first node; and 2) at least one recipient node beginning to communicate a portion of its respective subfile that it is receiving from the first node to at least one other recipient node before the at least one recipient node fully receives its respective subfile. The disclosure of each reference is addressed below to illustrate that the combination of the disclosures fails to teach or suggest at least this element of claim 1.

First, *Sim* fails to teach or suggest the above-emphasized elements of claim 1. In *Sim*, a large payload file is partitioned into a plurality of subfiles. The subfiles are all stored to an originating node. Certain ones of the subfiles are then distributed to other nodes for serving to end-user clients. However, *Sim* does not teach or suggest a distribution technique in which a first node attempts to distribute the plurality of subfiles that make up a file *F* from the first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of the first group but not all of the plurality of subfiles to any recipient node of the first group. Further, *Sim* fails to teach or suggest that the plurality of recipient nodes of the first group attempt to exchange their respective subfiles received from the first node. Rather, in *Sim* the subfiles desired by one recipient node are obtained from a nearest node possessing the desired subfiles. As discussed further below, no exchange of respective subfiles received from a first node is attempted between the recipient nodes of a first group in *Sim*.

*Sim* explains its system for distributing a large payload file at column 14, lines 10-42 as follows:

A content provider uploads a large payload file to a single content management server using content publishing and management tools running on a content provider client system. After receiving the file, the CMS processes the file and breaks it down, if required, into track files (a.k.a. linear files). A linear file comprises a file that maintains the order associated with the substance (i.e., substantive content) of the file. If, for example, the linear file contained a movie, the beginning of that file would include the beginning portions of the movie. Similarly, the middle and end portions of the movie would be located at the middle and end of the linear file. Linear files are desired because it is easier to reassemble such files using linear superposition, for example. Some media files are non-linear, that is, they contain multiple tracks such that the first part of the movie, for example, is not stored in the beginning of the file. After breaking the file down to linear (i.e., track) files, the CMS transfers the file to the distribution server it is connected to. The distribution server further breaks the track files down to block files, as desired for storage. The block files may subsequently be stored in local storage locations 711-713, for example. A file distribution protocol (e.g., FDP) command is subsequently used to distribute (i.e., replicate) the file, or selected portions thereof, to other distribution server nodes within the scalable content delivery network. For initial replication, the entire block files need not be stored in all nodes however a master copy may be maintained completely in one node (typically the originating node). The FDP includes commands to facilitate

file transfers and manipulations within the SCDN. The size of the blocks affects the performance of both content distribution and content delivery and is discussed later in this document.

*Sim* illustrates its distribution technique in FIG. 13, which it explains at col. 20, line 34 -- col. 21, line 24 as follows:

FIG. 13 is an illustrative embodiment of the distribution of a large payload file within an SCDN. A content provider uploads a large payload file into the content management server (CMS) 570, which is connected to node B of the SCDN, using any content publishing and management software running on the content provider's client system (CPC) 530. The content provider also uploads the distribution criteria onto CMS 570. Content management server 570, as previously described, divides the uploaded file into track files and issues a command similar to the FDP "put" command for each track file to the distribution server located in node B. In other embodiments, the CMS may be connected to any node of the SCDN. At node B, the DS divides the track files into block files for local storage. The full copy of the file is shown at Node B as a filled in dot. The CMS then issues an FDP command of the type "distribute" to the distribution server at node B. In response to the distribute command, the DS issues a command to its neighboring nodes A, D, and E to replicate the content (e.g., using the "replicate" command of the FDP). Node D examines the replicate packet and decides its not supposed to have the content thus it passes the replicate command to its neighbor, node H. Nodes A, E, and H examine the replicate packet and decide they all match the distribution criteria (i.e., they are "qualified" nodes). When ready, nodes A, E, and H issue commands to retrieve a portion of the file from the nearest node (e.g., node B) in the SCDN. Nodes E and H are leaf nodes thus they do not propagate the replicate command. However, node A is the root node with child nodes B and C. Node A may not send the replicate command back to node B, because it is the originating node. However, node A may send the replicate request to node C. Node C checks the distribution criteria and decides it's a qualified node therefore it retrieves a portion of the file from the nearest nodes (e.g., the nearest of nodes A, B, E, and H) containing the needed data. Node C subsequently sends the replicate command to nodes F and G. Node F is qualified thus it retrieves a portion of the file from the nearest nodes having the data (e.g. nodes B or C). Nodes G and H are not qualified thus they receive nothing. Node G is a terminating node because the rolled-up attribute of its branch does not satisfy the distribution criteria. This initial replication process continues until all the qualified nodes in SCDN are at least partially populated. In one or more embodiments, the same portion (e.g., blocks) of the large payload file is contained in at least one node of the SCDN. Preferably, a plurality of nodes maintains the same portion thereby creating redundancy and preventing loss of any portion of the large payload file when one or more nodes or storage volumes become

unavailable. For example, when a storage volume (or device) becomes unavailable (i.e., lost), a DS at that station need not take any special action to recover contents of the damaged volume since the portions of large payload files stored and hence lost in that volume are automatically downloaded from other network nodes upon demand to service a user request. The distribution servers also relay control information of a failed station to neighbors of the failed station to prevent improper termination of control commands.

*Sim* does not teach that the recipient nodes of a first group exchange their respective subfiles. Rather, as shown in FIG. 13, node B (an origin node) contains all of the subfiles. Node B sends a “distribute” command to its neighboring nodes A, D, and E, which is propagated through the other nodes H, C, and F-I, and such nodes determine based on a distribution criteria whether they are to obtain any of the subfiles. If they are to obtain any of the subfiles, the nodes A and C-I retrieve those subfiles that they are to obtain from a nearest node possessing such subfiles.

As shown in FIG. 13, nodes A, C, E, F, and H each receive some of the subfiles. However, *Sim* does not teach that any of the recipient nodes A, C, E, F, and H exchange their respective subfiles. While a given recipient node may receive its desired subfiles from another of the recipient nodes that is nearest it, *Sim* does not teach that the recipient nodes “exchange” their respective subfiles.

Further, in the example of FIG. 13 of *Sim*, assuming that nodes A, E, and H may be considered a “first group” which receive subfiles from origin node B, *Sim* provides no teaching whatsoever that the recipient nodes A, E, and H exchange their respective subfiles received from origin node B. Rather, in *Sim*, whatever subfiles are desired by recipient nodes A, E, and H, those recipient nodes obtain from a nearest node possessing the desired subfiles (e.g., origin node B). That is, *Sim* provides no teaching that nodes A, E, and H each receive subfiles from origin node B, and then attempt to exchange their respective subfiles received from node B. For example, in FIG. 13, it appears that nodes A and E each receive subfiles from node B. *Sim* provides no teaching that nodes A and E then exchange their respective subfiles received from node B with each other. Even in the event that node E later desires a subfile that is possessed by node A (and assuming that node A is nearest node E), node A may send such desired subfile to

node E, but *Sim* provides no teaching of an “exchange” in which node E also sends a subfile to node A.

Thus, *Sim* does not teach or suggest at least this “exchanging” element of claim 1.

Indeed, *Sim* teaches away from the recipient nodes exchanging their respective subfiles received from the origin node because *Sim* expressly attempts to avoid each recipient node receiving all of the subfiles. For instance, *Sim* teaches at col. 8, lines 59-62 that “[e]ach node at the edge of the network embodying aspects of the invention is configured to appear as if it has the large file stored locally when portions of the file are really stored on other nodes located throughout the network.” At col. 9, lines 2-4, *Sim* further explains that the “end result is that each network node has access to numerous large data files without having to store each of those data files locally.” As shown in the example of FIG. 13 of *Sim*, the origin node B contains all of the subfiles, while the recipient nodes A, C, E, F, and H each receive some of the subfiles, but not all of the subfiles. *Sim* does not provide any teaching or suggestion of the nodes exchanging their respective subfiles received from the first node.

The current Office Action appears to concede that *Sim* fails to teach or suggest the plurality of recipient nodes of the first group attempting to exchange their respective subfiles received from the first node, *see* Page 4 of the Office Action. However, the Office Action asserts that *Zayas* teaches this element of the claim. Applicant disagrees for the reasons discussed below.

*Zayas* is directed to an infrastructure in which volumes are replicated on each of a plurality of servers. The infrastructure permits different “replication modules” to be utilized for managing the replication (e.g., distributing updates to the volumes, etc.) of different volumes. For instance, FIG. 2A of *Zayas* shows an example in which a replication module A 265A is utilized for managing replication of a volume V3, while a different replication module B 265B is utilized for managing replication of volumes V1 and V2. Thus, *Zayas* explains that if one of the replication modules employed incurs large overhead for keeping the volumes that it manages

consistent, this does not lead necessarily to a large overhead for the volumes managed by a different replication module, *see e.g.* column 4, lines 26-35 of *Zayas*.

Thus, *Zayas* describes a system in which entire volumes of files are replicated onto a plurality of different servers. *Zayas* is not concerned, however, with how the files are distributed to the different servers. For instance, presumably one server in *Zayas* may simply distribute a full file to each of the other servers to which the file is to be replicated. *Zayas* is instead concerned with a system in which different modules can be used for managing updates to different volumes in order to maintain consistency in the volumes across the different servers to which they are stored. Accordingly, *Zayas* also fails to teach or suggest a plurality of recipient nodes attempting to exchange their respective subfiles received from a first node.

Further, the Office Action concedes on page 5 thereof that neither *Sim* nor *Zayas* teaches or suggests at least one recipient node beginning to communicate a portion of its respective subfile that it is receiving from the first node to at least one other recipient node before the at least one recipient node fully receives its respective subfile. However, the Office Action contends that *Chung* provides this element, citing to paragraphs 0008-0009 and 0016 of *Chung*. Applicant disagrees for the reasons discussed below.

While *Chung* mentions that subfiles may be distributed to different servers, and may then be streamed in parallel from the servers to an end user, *Chung* makes no mention whatsoever of a given recipient node of a subfile in a group beginning to communicate a portion of the subfile that the given recipient node is receiving to another node in the group before the given recipient node fully receives the subfile. First, the servers to which the subfiles are distributed do not exchange their respective subfiles, and *Chung* provides no teaching whatsoever of any of the servers beginning to communicate a subfile to another server before the subfile is fully received. Further, while subfiles may be streamed to an end user from the servers, *Chung* does not teach or suggest that a server begins streaming a subfile to an end user before the subfile is fully received by the server. Additionally, *Chung* provides no teaching that the end user communicates a received subfile to another recipient node, and certainly fails to teach or suggest that the end user



begins communicating the received subfile to another recipient node before the end user fully receives the subfile.

The Examiner appears to contend that the servers of *Chung* provide the recited first group comprising a plurality of recipient nodes. However, *Chung* provides no teaching of one server communicating a subfile that it receives to another of the servers. Further, if the Examiner further considers an end user of *Chung* as a recipient node in the recited first group, *Chung* likewise fails to teach or suggest that the end user communicates the subfile that it receives to another node. Thus, irrespective of how a given end user may receive a subfile from the servers (e.g., via streaming or other download technique), *Chung* provides no disclosure of the given end user communicating the subfile that it receives to another recipient nodes. *Chung* certainly provides no teaching of one recipient node (the end user or a server) beginning to communicate a portion of its respective subfile that it is receiving to another of the recipient nodes (another server or end user) before the one recipient node fully receives its respective subfile.

As mentioned above, the Final Office Action cites to paragraph 0016 of *Chung* as teaching this element of claim 1. Paragraph 0016 of *Chung* merely provides:

One aspect of the invention involves dividing a single file into multiple files or sub-files. A sub-file has a file name and other file attributes, and is treated by the operating system's file system as just another file. The divided files or sub-files may then be distributed and stored onto one or more servers. When an end user wants the file to be delivered in a streaming fashion, the sub-files can be transmitted in parallel and simultaneously from one or more servers, which increases the rate at which data can be delivered.

This portion of *Chung* mentions that the subfiles can be distributed (from an origination node that contains the full file) to the various servers, such as servers A-D of FIGURE 4. This portion of *Chung* further mentions that the subfiles from the servers can be distributed to an end user (client) in parallel and delivered to the end user in a streaming fashion. Thus, if a client such as client 610 in FIGURE 6 of *Chung* requests a file, the various subfiles S1-S4 may be sent to the client in parallel from the different servers A-D.

As discussed above, *Chung* provides no teaching or suggestion of a recipient node (either a server or an end user) communicating its subfile to another recipient node. Further, *Chung* provides no teaching or suggestion of one recipient node (either a server or an end user) beginning to communicate a portion of its respective subfile to another of the recipient nodes before the one recipient node fully receives its respective subfile. For instance, even when servers A-D communicate their respective subfiles to the client node 610, *Chung* makes no mention of the servers beginning such communication of a subfile from the server to the client before the subfile is fully received by the server.

Thus, *Chung* fails to teach or suggest the above element of claim 1. Neither *Sim* nor *Zayas* is relied upon as teaching such element, nor do they do so.

In view of the above, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 1, and thus the rejection of claim 1 should be overturned.

#### Independent Claim 17

Independent claim 17 recites:

A system comprising:  
an origin node operable to partition a file *F* into a plurality of subfiles, wherein said plurality of subfiles correspond in number to a number of recipient nodes in a first group to which said file is to be distributed;  
said origin node operable to attempt to distribute all of said plurality of subfiles to said recipient nodes, wherein said origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes;  
said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node;  
said origin node operable to detect a failed node in said first group; and  
said origin node operable to manage distribution of said file *F* upon detecting a failed node in said first group in a manner such that every non-failed node of said first group receives said file *F*. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 17. As discussed above with claim 1, the applied combination fails to teach or suggest at least: 1) said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node, and 2) at least one recipient node that begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node. Further, the applied combination fails to teach or suggest such an attempt to exchange the subfiles “such that each recipient node obtains all of said plurality of subfiles”. Indeed, as discussed further herein, *Sim* and *Chung* each expressly teach away from attempting to have each recipient node obtain all of the plurality of subfiles.

Thus, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 17, and thus the rejection of claim 17 should be overturned.

#### Independent Claim 21

Independent claim 21 recites:

A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

attempting to distribute a plurality of subfiles that comprise a file *F* from a first node to a first group comprising a plurality of recipient nodes, wherein the first node attempts to distribute at least one subfile to each recipient node of said first group but not all of said plurality of subfiles are distributed from the first node to any of the recipient nodes of said first group;

said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting whether one of said plurality of recipient nodes of said first group has failed; and

if a recipient node of said first group has failed, managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that the file *F* is distributed to each non-failed node of said plurality of recipient nodes. (Emphasis added).

The combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest at least the above-emphasized element of claim 21. As discussed above with claim 1, the applied combination fails to teach or suggest at least: “said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile”.

Thus, the applied combination of *Sim*, *Zayas*, and *Chung* fails to teach or suggest all elements of claim 21, and thus the rejection of claim 21 should be overturned.

#### Dependent Claims

Claims 2-16, 18-20, and 22-28 each depend from one of independent claims 1, 17, and 21, and are thus likewise believed to be allowable at least based on their dependency from their respective independent claim for the reasons discussed above. Accordingly, Applicant respectfully requests that the rejection of claims 2-16, 18-20, and 22-2 also be withdrawn.

#### *2. Insufficient Motivation to Combine the References in the Manner Applied*

Further, insufficient motivation exists for combining the teachings of *Sim*, *Zayas*, and *Chung* in the manner relied upon by the Final Office Action. The Final Office Action alleges that one of ordinary skill in the art would be motivated to combine *Sim* and *Zayas* to arrive at a system in which a plurality of nodes of a first group attempt to exchange their respective subfiles received from a first node (or origin node), *see e.g.*, the rejection of claim 1. However, *Sim* expressly teaches away from any such system that attempts a distribution that results in each recipient node obtaining all of the subfiles, as discussed above with claim 1. For instance, *Sim* teaches at col. 8, lines 59-62 that “[e]ach node at the edge of the network embodying aspects of the invention is configured to appear as if it has the large file stored locally when portions of the file are really stored on other nodes located throughout the network.” At col. 9, lines 2-4, *Sim* further explains that the “end result is that each network node has access to numerous large data

files without having to store each of those data files locally.” Thus, *Sim* expressly attempts to avoid each network node from storing the full file locally.

On the other hand, *Zayas* is directed to a system in which full volumes of files are replicated on different servers and replication modules are employed for maintaining consistency among the volumes on the various servers. Thus, *Zayas* goes directly against the express teaching/desires of *Sim*. As such, one of ordinary skill in the art would not be motivated to combine the teaching of *Zayas* with the teaching of *Sim*. Thus, for at least this reason, the rejections should be withdrawn.

Further, independent claim 17 recites, in part, “said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles” (emphasis added). However, *Chung* expressly teaches away from any such system that results in each recipient node attempting to obtain all of the subfiles. *Chung* expressly teaches at paragraphs 0005-0006 that replicating a full file onto a large number of servers is undesirable because it uses large amounts of expensive disk storage, etc. Thus, *Chung* proposes that a file be divided into a plurality of subfiles that are distributed to different servers without requiring that the entire file be distributed to each server.

On the other hand, *Zayas* is directed to a system in which full volumes of files are replicated on different servers and replication modules are employed for maintaining consistency among the volumes on the various servers. Thus, *Zayas* goes directly against the teaching of *Chung*. As such, one of ordinary skill in the art would not be motivated to combine the teaching of *Zayas* with the teaching of *Chung*. Thus, for this further reason, insufficient motivation for combining *Chung* and *Zayas* exists, and thus the rejections should be withdrawn.

**V. Conclusion**

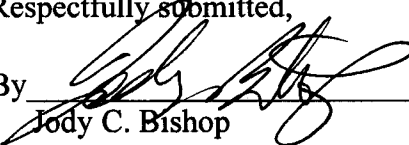
In view of the above, Applicant believes the pending application is in condition for allowance.

The required fee for this response is enclosed. If any additional fee is due, please charge Deposit Account No. 08-2025 under Order No. 200310236-1, from which the undersigned is authorized to draw.

Dated: December 19, 2006

Respectfully submitted,

By



Jody C. Bishop

Registration No.: 44,034

FULBRIGHT & JAWORSKI L.L.P.

2200 Ross Avenue, Suite 2800

Dallas, Texas 75201-2784

(214) 855-8000

(214) 855-8200 (Fax)

Attorney for Applicant